

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**System, Method, and Computer Program Product for  
File Encryption, Decryption and Transfer**

Inventors:

David Cross

Robert Gu

Duncan Bryce

Shishir Pardikar

Pradeep Madhavarapu

Scott Field

Kelvin Yiu

ATTORNEY'S DOCKET NO. MS1-1627US

CLIENT'S DOCKET NO. 305131.1

**EL9962768 30**

# **SYSTEM, METHOD, AND COMPUTER PROGRAM PRODUCT FOR FILE ENCRYPTION, DECRYPTION AND TRANSFER**

## **TECHNICAL FIELD**

[0001] The described subject matter relates to digital computing devices, and particularly to systems, methods, and computer program products for encrypting and decrypting files, and for transferring encrypted files between computing devices.

## **BACKGROUND**

[0002] The protection of sensitive data has become an important issue to users of computers. Sensitive data, such as corporate records, credit card numbers, personal data, or other sensitive information may be misappropriated if a computer is stolen, or if an unauthorized individual gains access to a computer network and copies information from some or all of its files or by monitoring and recording the network packets that are transmitted unprotected. Those authorized to access the sensitive information may not even know that the information has been copied.

[0003] To protect information, one type of security procedure involves encrypting the data, so that even if the data falls into the wrong hands, it cannot be read by a party unless the party obtains a “key” to decrypt the data.

[0004] In a distributed computing environment, data may be transmitted between computers, such as, e.g., a client computer and one or more server computers, or between peer computers in a peer-peer environment. To reduce the

likelihood that an unauthorized party might obtain access to data passing between computers, it is desirable that encryption operations be controlled from the client so that users are in control of the encryption status of their files regardless of where the files are located. In a client-server environment, it is desirable that most encryption operations take place primarily on the client so that the server components or the server administrator(s) need not be a trusted component or a component that could intercept or subvert the system.

### **SUMMARY**

[0005] Implementations permit users of client computers to control encryption operations of files, even if the files are resident on a server. An encrypting file system (EFS) and an underlying file transfer protocol to permit a client to encrypt or decrypt a file resident on a server. In addition, a user at a client computer can open, read, and write to encrypted files, including header information associated with encrypted files, and can add users or remove users from an encrypted file.

[0006] In one exemplary implementation, a method of computing is provided. The method may be implemented at a server, which receives a request from a client to encrypt a first file resident on the first computing device. In response to the request, the server creates a second file, copies metadata from the

first file to the second file, and transmits an identifier associated with the second file to the client. The client may then complete the encryption process

[0007] In another exemplary implementation, a method of computing is provided. The method may be implemented at a client computer, which receives a request to encrypt a file from an application. The method comprises generating a request to encrypt a first file, transmitting the request to a second computing device, which may be a server; receiving, from the second computing device, an identifier associated with a second file; opening the first and second files; writing header information into the second file; encrypting the contents of the first file; writing the encrypted contents into the second file; and replacing the first file with the second file.

[0008] In yet another exemplary implementation, a method of computing is provided. The method comprises generating, at a client computing device, a request to encrypt a first file and transmitting the request to a server computing device. In response to the request, the server creates a second file; copies metadata from the first file to the second file, and transmits an identifier associated with the second file to the client computing device. In response to the receipt of the identifier, the client opens the first and second files, writes header information into the second file, encrypts the contents of the first file, writes the encrypted contents into the second file, and replacing the first file with the second file.

[0009] In another exemplary implementation, a computer system is provided. The computer system comprises a display, a user-input device, and a processor capable of executing logic instructions. The computer system further comprises a computer readable medium comprising logic instructions for receiving a request to encrypt a first file resident on the first computing device, and in response to the request, creating a second file, copying metadata from the first file to the second file, and transmitting an identifier associated with the second file to the first computing device.

[0010] In another exemplary implementation, a method of computing is provided. The method comprises receiving, at a first computing device, a request from a second computing device to decrypt a first file resident on the first computing device. In response to the request, the first computing device creates a second file, copies metadata from the first file to the second file, decrypts the first file, writes contents of the first file to the second file, and replaces the first file with the second file. The first computing device may be implemented as a server and the second computing device may be a client.

[0011] In another exemplary implementation, a method of computing is provided. The method comprises generating, at a first computing device, a request to decrypt a first file, obtaining the File Encryption Key (FEK) of the first file, and forwarding the FEK to a second computing device with the request to decrypt the

first file. The first computing device may be implemented as a client and the second computing device may be a server.

[0012] In another exemplary implementation, a method of computing is provided. The method comprises generating, at a first computing device, a request to decrypt a first file, obtaining the FEK of the first file; and forwarding the FEK to a second computing device with the request to decrypt the first file. In response to the request, the second computing device creates a second file, copies metadata from the first file to the second file, decrypts the first file, writes contents of the first file to the second file, and replaces the first file with the second file. The first computing device may be implemented as a client, and the second computing device may be a server.

[0013] In another exemplary implementation, a computer system is provided. The system comprises a display, a user-input device, a processor capable of executing logic instructions. The system further comprises a computer readable medium comprising logic instructions for receiving a request from a second computing device to decrypt a first file resident on the first computing device, creating a second file, copying metadata from the first file to the second file, decrypting the first file, writing contents of the first file to the second file, and replacing the first file with the second file.

[0014] In another exemplary implementation, a method of adding a user to an encrypted file stored on a server is provided. The method comprises obtaining

the \$EFS stream from a server, generating a data decryption field (DDF) for the user, modifying the \$EFS stream to include the DDF for the user, and writing the \$EFS stream to the server.

[0015] In another exemplary implementation, a method of removing a user from an encrypted file stored on a server is provided. The method comprises obtaining the \$EFS stream from a server, locating a DDF corresponding to the user in the \$EFS stream, modifying the \$EFS stream to include the DDF for the user, and writing the \$EFS stream to the server.

### **BRIEF DESCRIPTION OF THE DRAWINGS**

[0016] Fig. 1 is a block diagram representing an exemplary computer system;

[0017] Fig. 2 is a block diagram representing an exemplary component architecture;

[0018] Fig. 3 is a block diagram conceptually representing various logical components used in the encryption of data;

[0019] Fig. 4 is a block diagram conceptually representing various logical components used in the decryption of data;

[0020] Fig. 5 is a block diagram conceptually representing various logical components used in the recovery of encrypted data;

[0021] Fig. 6 is a representation of stream control blocks associated with files, at least some of which include key context information for file encryption and decryption;

[0022] Fig. 7 is a representation of a context chain used for communicating information between encryption components for encrypted files and directories;

[0023] Figs. 8 and 9 are representations of data structures used by certain components for communicating file information, including encryption key information, to one another;

[0024] Fig. 10 is a flow diagram illustrating an exemplary flow of control to open or create a file;

[0025] Fig. 11 is a flow diagram representing preprocessing steps generally taken as part of opening or creating a file;

[0026] Fig. 12 is a flow diagram representing steps taken by the file system to handle a file open request;

[0027] Figs. 13 – 14 comprise a flow diagram representing exemplary steps taken by a callout to open or create an encrypted file;

[0028] Figs. 15 – 19 comprise a flow diagram representing post-processing steps generally taken as part of opening or creating a file;

[0029] Fig. 20 is a flow diagram representing steps taken by the various components to handle a file read request;



[0030] Fig. 21 is a flow diagram representing steps taken by the various components to handle a file write request;

[0031] Fig. 22 is a flow diagram representing steps taken by the various components to handle a user request to encrypt a stored plaintext file;

[0032] Fig. 23 is a flow diagram representing steps taken by the various components to handle a user request to decrypt a stored encrypted file;

[0033] Fig. 24 is a block diagram representing an exemplary component architecture for a client-server EFS system;

[0034] Fig. 25 is a flow diagram illustrating an exemplary read operation;

[0035] Fig. 26 is a flow diagram illustrating an exemplary write operation;

[0036] Fig. 27 is a flow diagram illustrating an exemplary encrypt operation;

[0037] Fig. 28 is a flow diagram illustrating an exemplary decrypt operation;

[0038] Fig. 29 is a flow diagram illustrating an exemplary method for opening an encrypted raw file in a client-server context;

[0039] Fig. 30 is a flow diagram illustrating an exemplary method for opening an encrypted raw file in a client-server context;

[0040] Fig. 31 is a flow diagram illustrating an exemplary procedure for deleting a user from an encrypted file; and

[0041] Fig. 32 is a flow diagram illustrating exemplary operations for creating a file.

## **DETAILED DESCRIPTION**

### Overview

[0042] Described herein are exemplary systems and methods for encrypting, decrypting, and transferring files between computing devices. The methods described herein may be implemented as logic instructions on a computer-readable medium and executed on a processor to configure computing devices to function as a special purpose machines for encrypting, decrypting, and transferring files.

[0043] In exemplary embodiments, the systems and methods described herein may be implemented in a client-server computer network in which the computer processors execute the WINDOWS® brand operating system commercially available from Microsoft Corporation of Redmond, Washington, USA. Exemplary systems and methods described herein invoke the Encrypting File System (EFS) and the NT File System (NTFS) functionality provided in the WINDOWS® brand operating system to facilitate encrypting, decrypting, and transferring files. It will be appreciated, however, that the systems and methods described herein are not limited to the WINDOWS® brand operating system. Rather, these systems and methods may be implemented on other operating systems including, but not limited to UNIX operating systems and its variants.

[0044] Broadly, the systems and methods described herein permit encryption and decryption of computer files to be controlled and implemented from a client computer, even if the target computer files reside on a server computer (or another client computer in a peer-to-peer network). Protocols are implemented at both the client and the server to permit control of encryption at the client.

[0045] Following is a description of an exemplary EFS and NTFS that may be invoked to implement the systems and method described herein. The exemplary EFS and NTFS are described in U.S. Patent Application Pub. No. 2002/0019935 to Andrew, et al., filed May 29, 2001, the entire disclosure of which is incorporated by reference herein.

#### Exemplary Computer System

[0046] Turning to the drawings and referring first to **Fig. 1**, there is shown an exemplary computer system 20 on which systems and methods described herein may be implemented. The computer system 20 may be a server, a workstation, or a combination thereof, and may be connected in a known manner to one or more other computer-based resources.

[0047] As shown in FIG. 1, computer system 20 includes a processor 22 connected to a memory 24 having an operating system 26 loaded therein. One suitable operating system 26 is Microsoft Corporation's Windows® 2000 operating system. The computer 20 has a file system 28 such as the Windows NT File

system 28 (NTFS 28) associated with or included within the operating system 26. However, as can be appreciated, the present invention is not limited to any particular operating system and/or file system, but for clarity the present invention will be hereinafter described with reference to Windows® 2000 and NTFS 28. At least one application program 30 in the memory 24 interfaces with the operating system 26 and the file system 28 through application programming interfaces (APIs) 32.

[0048] The computer system 20 also includes input-output (I/O) circuitry 34 for connecting the computer system 20 to one or more networked devices, to one or more input devices 36 such as a keyboard and/or mouse, and/or to one or more output devices 38 such as a monitor and/or speakers. The computer system 20 also includes a non-volatile storage device 40 such as a hard disk drive. As can be appreciated, the non-volatile storage device 40 may also operate in conjunction with the random access memory of the memory 24 to provide a large amount of virtual memory via swapping techniques.

[0049] The file system 28 connects through a device driver 42 to communicate with the non-volatile storage device 40, to manage the files thereon, and generally contains methods for (1) storing, referencing, sharing and securing files, (2) accessing file data and (3) maintaining file integrity. Notwithstanding, there is not always a clear distinction between a file system 28 and its associated operating system, particularly with those file systems 28 contained within an

operating system. Accordingly, it is understood that any or all of the processes or steps attributed herein to the file system 28 may alternatively be performed by the operating system 26, and vice-versa.

[0050] The non-volatile storage 40 stores a number of files 441 - 44n, which, when NTFS serves as the file system 28, have their data organized in attribute data streams. An NTFS file control block (FCB) associated with each file maintains information identifying the data streams belonging thereto. Windows NT and NTFS 28 are described in the texts, *Inside Windows NT*, by Helen Custer, Microsoft Press (1993) and *Inside the Windows NT File System*, Helen Custer, Microsoft Press (1994). As shown in FIG. 1 and as described below, in accordance with the present invention, at least some of the files 441 - 44n are stored with encrypted data.

#### EFS Component

[0051] Referring to **Fig. 2**, an encrypting file system is provided to encrypt and decrypt files. In an exemplary embodiment, the encrypting file system comprises an Encrypting File System (EFS) linked library 47, (*e.g.*, DLL), an EFS runtime library (FSRTL 48) 48 and an EFS service 50. The linked library 47 provides relatively tight integration with the file system, (as opposed to an installable filter driver model such as described in United States Patent Application Serial No. 08/931,774, the entire disclosure of which is incorporated herein by

reference), *e.g.*, both are loaded together (and not according to a registry setting). Notwithstanding, the EFS may be implemented with a linked library, driver, or virtually any other mechanism including directly incorporating the encryption functions into the file system code.

[0052] The EFS linked library 47 registers with the file system 28, whereby the file system provides encryption functionality that is transparent (*e.g.*, to an application) by calling the EFS linked library's functions, listed in a function table 47A or the like acquired by the file system during registration. Note that instead of linking in this manner, these functions may be incorporated into the file system 28, however the modularity of these components provides benefits normally associated with modularity. For purposes of this description, once registered, the EFS linked library 47 generally may be considered part of the file system 28. Further, note that if for some reason the EFS linked library 47 cannot link to and register with the file system, (*e.g.*, errors may occur during the initialization phase), then the file system will not provide encryption and decryption functionality. For example, a user will not be able to access an encrypted file (until the library is properly initialized).

[0053] During initialization, the EFS linked library 47 registers file system runtime library callback routines (FSRTL 48 routines) with the NTFS 28, maintained in the function table 47A. As described below, NTFS 28 uses these FSRTL 48 routines to call back to obtain file encryption related services.

[0054] The EFS linked library 47 provides the support to communicate with the user mode EFS service 50 running as part of the security subsystem. During initialization (or alternatively when encryption or decryption is first needed), the EFS linked library 47 communicates with the EFS service 50 using a GenerateSessionKey interface, to establish a symmetric session key that is used to communicate securely between the EFS linked library 47 and the EFS service 50. Data communicated between the two is encrypted using this session key. This session key is also used by callouts to the FSRTL 48 to decrypt I/O controls from the EFS service 50.

[0055] During open of an encrypted file, the EFS linked library 47 communicates with the EFS service 50 by passing it the file metadata, including the data decryption and data recovery fields, (Figs. 3–5, described below), to get back the file encryption key and any updates to the file metadata. The file metadata may be updated because the user may change to a new key, or the recovery agent's keys might get updated. The EFS linked library 47 passes this information to FSRTL 48.

[0056] During encryption of a plaintext file/directory or creation of a new encrypted file, the EFS linked library 47 communicates with the EFS service 50 to get a new file encryption key, and encryption metadata for the encrypted file. The EFS linked library 47 also passes this information to the FSRTL 48.

## EFS FSRTL

[0057] The FSRTL 48 is a module that implements NTFS callouts to handle various file system 28 operations such as reads, writes, and opens, on encrypted files and directories, as well as operations to encrypt, decrypt, and recover file data when it is written to or read from disk. To this end, the present invention provides a callout mechanism including an interface between NTFS 28 and the FSRTL 48. As described in more detail below, this interface is generic to any appropriate library (and driver) that transform data, including the ones described herein that encrypt data, and thus the interface between NTFS 28 and FSRTL 48 is more accurately referred to as a data transformation interface 52. For example, an indexing driver could use this interface to monitor all writes to disk and develop an index based on those writes. However, as can be appreciated, a dedicated encryption interface may be alternatively provided. Note that in one preferred implementation the interface 52 (and other interfaces herein) is generic to allow for other EFS-like packages to be supplied).

[0058] Operations between the EFS linked library 47 and FSRTL 48 include writing EFS attribute data (decryption data and recovery fields) as file attributes, and communicating a file encryption key computed in the EFS service 50 to FSRTL 48, such that it can be set up in the context of an open file. This file context is then used for transparent encryption and decryption on writes and reads of file data to and from the non-volatile storage 24.



[0059] The data transformation interface 52 is capable of interfacing to any engine or driver that transforms the data in virtually any way, but for purposes herein the interface 52 will be described as interfacing the EFS linked library 47 to the file system 28 for accomplishing data encryption. Notwithstanding, the data transformation interface is not limited to data encryption, but is appropriate for accomplishing virtually any type of data alteration. At present, however, this transformation model supports in-place data transformation wherein the data takes at least no more space than the original plain text. In any event, the EFS linked library 47 registers these callbacks with the file system 28, whereby the file system 28 uses the registered EFS callback functions at appropriate times to carry out the various encrypting and decrypting tasks that the user requests.

[0060] Although not necessary to the invention, for convenience, the FSRTL 48 is stored in a common file with the EFS linked library 47. Indeed, although the EFS linked library 47 and FSRTL 48 are implemented as a single component, they do not communicate directly, but instead use the NTFS file control callout mechanism, *i.e.*, the EFS linked library 47 can effectively call the FSRTL 48. The use of the NTFS callout mechanism ensures that NTFS 28 participates in all file operations, which avoids conditions such as where two users are locked, each user waiting for the release of the other's file.

[0061] The data transformation interface 52 includes a number of function pointers, or callbacks. A first callback which the file system 28 uses, the

FileCreate callback, tells the registered EFS functions that a stream is being created or opened. The actions that EFS linked library 47 takes at this point (*e.g.*, determining if a user has access if the file is an existing file or getting the metadata stream for a new file) are described in more detail below.

[0062] When an application opens or creates a file, the I/O subsystem 56 determines the file is of a certain file system, *e.g.*, an NTFS 28 file, and passes the request on to NTFS 28. NTFS 28 determines whether EFS may be interested in the file, *e.g.*, if the file is created in an encrypted directory or if a stream is created or attached to an encrypted file. IF NTFS 28 determines that the file is of interest to EFS, and sees that the EFS linked library 47 is registered therewith, NTFS 28 calls a registered EFS function, *i.e.*, the FileCreate callback. If the request is a file open request on an existing file, FSRTL 48 reads the file metadata from the file attribute and fills up a context block (*e.g.*, block 981 of FIG. 7, previously allocated by the EFS linked library 47, as described below) to pass back that information to the EFS linked library 47. When the call returns from NTFS 28, the EFS linked library 47 takes the metadata information and communicates with the EFS service 50 to extract the file encryption key 60 from the metadata. This information is then returned by the EFS linked library 47 to NTFS 28 by another FSRTL 48 interface, FileControl, described below, which sets up a key context 96 on the file being opened. This key context 96 is thereafter retained by NTFS 28 for future calls to the EFS linked library 47 until the file is closed. If the file

metadata is updated, the updated metadata is also re-written to the attributes by the registered EFS functions through NTFS callbacks.

[0063] If a new file is created, the FileCreate call results in the FSRTL 48 filling up the context buffer 981 with a request for a new file encryption key and metadata. The FSRTL 48 then passes the context buffer 981 back to the EFS linked library 47. The EFS linked library 47 takes this information and communicates with the EFS service 50 to obtain a new file encryption key and new file metadata from the EFS service 50. Using a file control callback (described below), the EFS linked library 47 returns this information to the FSRTL 48, whereby, using NtOfs function calls, the FSRTL 48 sets up the key context 98 on the file being created and writes the file metadata. The NtOfs API is a set of NTFS 28 function calls that allow the EFS linked library 47 to call into the file system 28 to manipulate the data streams containing the encryption meta data.

[0064] Another callback, FileSystemControl\_1, is called by NTFS 28 in response to the EFS linked library 47 request when a user is setting the encryption state of a file (EFS\_SET\_ENCRYPT), either marking it as encrypted or decrypted. In response, NTFS 28 sets or clears the encryption bit, and the EFS linked library 47 generates any necessary key storage. EFS\_SET\_ENCRYPT also originates in the EFS service 50 when a plaintext file begins to be encrypted, whereby the file state is modified such that no other operations are allowed on the file until the encryption is completed.

[0065] NTFS 28 also calls the FileSystemControl\_2 interface with various encryption driver-specific file control requests from the EFS linked library 47. Note that NTFS 28 takes no action with these callbacks other than to simply pass the call to the FSRTL 48. The file control requests include EFS\_SET\_ATTRIBUTE, which comes from the EFS filter EFS linked library 47 when it wants to write new or updated file metadata, and EFS\_GET\_ATTRIBUTE, which may come from the EFS linked library 47 or a user mode application 30 to query the file metadata. The information includes the list of user public keys and recovery agent public keys (described below) that are used to encrypt the file encryption key. Another request, EFS\_DECRYPT\_BEGIN, comes from the EFS service 50 when it starts decrypting an encrypted file. In response, the state of the file is modified such that no other operations are allowed on the file until the decryption is completed. EFS\_DEL\_ATTRIBUTE is a request originating in the EFS service 50 when it finishes decrypting an entire encrypted file, and wants to delete the file metadata and associated attribute. The EFS\_ENCRYPT\_DONE request also comes from the EFS service 50 when it successfully completes the file encryption. The file state is modified to allow any operations from this point on. EFS\_OVERWRITE\_ATTRIBUTE comes from the EFS service 50 when an encryption file is restored from its backup format. The EFS service 50 supplies the file metadata that needs to overwrite any existing metadata on the file. This

request is also associated with the deletion of any key context 96 associated with that file, such that no reads or writes can proceed while the file is being restored.

[0066] The FileSystemControl\_2 interface is also called by the file system 28 in response to the FSCTL\_ENCRYPTION\_FSCTL\_IO, also described below. This provides a means for the EFS linked library 47 to have NTFS 28 call the EFS linked library 47 (itself), such as when NTFS 28 recognizes that a file is in a certain state corresponding to a state for which the EFS linked library 47 is waiting.

[0067] The file system 28 directly uses the callback, AfterReadProcess after it has read some data from the disk for an encrypted file, and before returning it to the user. The AfterReadProcess function decrypts the data on the fly in response to this callback. The read operation is described in more detail below with respect to FIG. 20.

[0068] Conversely, BeforeWriteProcess is called by the file system 28 before it writes some data to the disk for an encrypted file. The function encrypts the data as a result of this callback. The write operation is described in more detail below with respect to FIG. 21.

[0069] The CleanUp callback is called by the file system 28 when NTFS 28 is freeing any resources associated with a stream. At this time, the EFS linked library 47 frees up any memory resources it was using, such as to store keys and the like. When NTFS 28 receives its last close on a stream, NTFS 28 performs its

normal operations to free up anything stored in memory to keep track of this open file, including the key context 96. In addition, the file system 28 calls the EFS linked library 47 with the context block 98, giving it the opportunity to free up any memory it was consuming for this file, *e.g.*, the context block 98.

[0070] The AttachVolume callback is called by a file system 28 (during the first user operation involving encryption), as described above. In response, the EFS linked library 47 notifies the I/O subsystem that it wants to attach to the device object representing that volume, thereby logically placing itself above NTFS 28 for that volume so that the I/O subsystem will pass information to the EFS linked library 47 first. DismountVolume is called by a file system 28 when a volume is being dismounted, either because a user wishes to eject or remove the drive, or because the system is being shut down. In response to the DismountVolume call, an encryption library or driver may free any memory resources that were allocated during the AttachVolume callback. However, it should be noted that the EFS linked library 47 ordinarily detaches itself and frees any resources when notified by the I/O subsystem of a volume dismount, but the DismountVolume callback is provided anyway to provide additional flexibility.

## EFS Service

[0071] The EFS service 50 is part of the Windows NT security subsystem. Referring to **Fig. 2** as EFS service 50 / Driver Communication 54, the EFS service

50 uses the existing local procedure call communication port between a Local Security Authority (LSA) and the kernel mode security reference monitor to communicate with the EFS linked library 47. In the user mode, the EFS service 50 interfaces with Microsoft's Cryptography API, CryptoAPI 58, to provide file encryption keys, and generate decryption field information.

[0072] The EFS service 50 also provides support for Win32 APIs 32, which are programming interfaces for encrypt, decrypt, recover and provide support for importing and exporting encrypted files. Importing and exporting encrypted files allows users to convert the files into opaque data (encrypted) for operations such as backup, restore, and general file transfer purposes as described below. The Win32 APIs 32 provide programming interfaces for encrypting plain text files, decrypting or recovering ciphertext files, and importing and exporting encrypted files (without decrypting them first). These APIs 32 are supported in a standard system DLL, advapi32.dll.

[0073] The EFS service 50 provides a number of services, including generating a session key and exchanging it with the EFS linked library 47 and the FSRTL 48. Based on the EFS linked library 47 request, the EFS service 50 generates a cryptographically strong session key (using CryptoAPI) and communicates it to the driver and FSRTL 48. The EFS service 50 also generates the file encryption keys in fields stored with the file (the Data Decryption Field, or DDF, and the Data Recovery Field, or DRF, described below with reference to

Figs. 3–5) using the user’s and recovery agents’ public key defined for EFS. When the EFS linked library 47 requests a new file encryption key, the EFS service 50 generates this information using CryptoAPI and returns it to the EFS linked library 47.

[0074] The EFS service 50 also extracts the file encryption key, *i.e.*, when the EFS linked library 47 requests this operation, the EFS linked library 47 supplies the file metadata, including the DDF and DRF key fields. Based on that information, the EFS service 50 sequentially searches the DDF and (if necessary) the DRF key fields to extract the name of the user’s key therefrom, and accesses its private portion via the CryptoAPI provider 58. If successful, (as described in more detail below), it passes the encrypted file encryption key to the provider for decryption. The service verifies that the decryption was correct (as also described below), and also verifies that the keys used in the file metadata are up to date. If the keys are not up to date, the service regenerates the metadata (DDF and/or DRF) and returns the extracted file encryption key and the metadata back to the EFS linked library 47.

[0075] When the EFS linked library 47 is loaded by NTFS 28, it first initializes its structures, and reserves some space to ensure that some memory is always available thereto. Then, the EFS linked library 47 registers itself with NTFS 28. Lastly, to synchronize with the driver, the EFS linked library 47 attempts to create a new event. If the event is successfully created, this indicates



that the EFS service 50 has not been initialized and the EFS linked library 47 has been loaded first. If successful, the EFS linked library 47 then creates a thread waiting on the event to be signaled. Later, when the event is signaled, *i.e.*, the EFS service 50 is ready to communicate, the EFS linked library 47 calls the EFS service 50 to get the session key. Once the session key has been transferred from the EFS service 50 to the EFS linked library 47, and the EFS service 50 and the EFS linked library 47 are synchronized. Note that if the event was not successfully created, it is ordinarily because the EFS service 50 was already initialized, in which event the EFS linked library 47 simply makes the call to get the session key.

[0076] In the situation where the EFS service 50 was loaded first, the EFS service 50 tries to create a new event. If the event is successfully created, then the EFS linked library 47 has not been initialized. The EFS service 50 generates the session key without waiting on the event. Later, when the EFS linked library 47 is loaded, the EFS service 50 will be called by the EFS linked library 47 to provide the session key thereto. When the EFS service 50 provides the session key, the EFS service 50 closes the event which was created earlier by the EFS service 50, and the EFS service 50 and the EFS linked library 47 are synchronized. Note that if the event was not successfully created, it is ordinarily because the EFS linked library 47 was already initialized, in which event the EFS service 50 instead opens the event and signals the event to let the EFS linked library 47 know that the EFS

service 50 is ready. Thereafter the EFS service 50 is asked for the session key by the EFS linked library 47, and the synchronization is done.

## System APIs

[0077] As described in more detail below with particular respect to **Figs. 22-23**, the EFS service 50 also provides a number of other user mode interfaces. These interfaces work closely with system APIs (*e.g.*, WIN32 or the like) to enable the user to perform operations such as convert an existing plaintext file to an encrypted file, convert an encrypted file to a plaintext file, and provide a backup and restore mechanism. By way of example, the Win32 interfaces 32 work with the EFS service 50 to expose EFS functionality, and include **EncryptFile**, which is a wrapper that calls into the interface provided by the EFS service 50 to do file encryption (FIG. 22). Another interface, **DecryptFile**, is a wrapper that similarly calls into the interface provided by EFS service 50 to do file decryption/recovery (FIG. 23).

[0078] A Backup/Restore mechanism is also provided in the system APIs, which enables users and backup operators to backup encrypted files without decryption. To this end, an OpenRawFile interface allows the user to open an encrypted file without read access, and without setting up a file encryption key to do transparent reads and writes. For these operations, NTFS 28 recognizes the access level and does not call the encryption EFS linked library 47 to look up a key

for this file, nor to decrypt reads nor encrypt writes. The only operations allowed on a file opened via this interface are file controls. Thus, a ReadRawFile interface allows the user to read all the data from the file, including the encryption metadata, as a contiguous opaque stream that can be backed up and later restored. A WriteRawFile interface allows the user to write all the data to the file from its backup, including the encryption metadata, to re-create the encrypted file. Lastly, a CloseRawFile is provided that allows the user to close the file which was opened raw by OpenRawFile.

[0079] A FileControl interface allows the Win32 APIs that provide the backup and restore mechanism to read and write raw encrypted data. Note that such raw data reads and writes are from/to NTFS 28 direct to/from the disk (storage); EFS becomes involved because the EFS service 50 and the EFS linked library 47 share the common session key, and all file controls need to be verified (as described below). For backing up the file, the Win32 APIs read the EFS metadata via an EFS file control, which translates into the FileSystemControl\_2 that returns the EFS stream. Then, NTFS 28 file controls are called to read the actual file data, which is packaged into an opaque stream and written out. To (gradually) restore the file, the reverse process is performed. An EFS file control is called to identify a first stream and another EFS file control to write the raw data back.

## Data Encryption

[0080] As conceptually represented in **Figs. 3-5**, data encryption and decryption may be implemented using a public key-based scheme. To this end, file data is encrypted using a fast symmetric algorithm with a file encryption key (FEK) 60 (FIG. 3). The FEK 60 is a randomly generated key of a certain length required by the selected algorithm, or as otherwise required if the algorithm supports variable length keys. As represented in Fig. 3, a random number generator 62 generates the FEK 60. To encrypt the file data using the FEK 60, the plain text 64 of the file is encrypted by a file encryption mechanism 66 using an appropriate algorithm (*e.g.*, DES, AES, or another encryption algorithm) and written as encrypted text 68 to an encrypted file 70.

[0081] As shown in Fig. 3, the randomly generated FEK 60 is itself encrypted with the public key 72 of at least one user, and stored with the encrypted file 70 in a special EFS attribute called the Data Decryption Field (DDF) 74. Using a suitable encryption algorithm, (*e.g.*, RSA), a data decryption field generator 76 performs the key encryption. In keeping with public-key based schemes, the private portion of the user's key pair is only used during decryption, *i.e.*, an encrypted FEK 60 in the data decryption field 74 is decrypted using the private portion of the key pair. The private portion 84 (Fig. 4) of a user key pair is safely stored in a separate location, such as on a smart card and/or other secure storage device. Note that encryption can also be done using a symmetric

algorithm, such as a password-derived key, but while feasible, EFS preferably does not support such encryption because password-based schemes are inherently weak due to dictionary attacks and the like.

[0082] As also represented in Fig. 3, the FEK 60 is also encrypted using one or more recovery public keys 78. The recovery key encryption public keys 78 belong to trusted persons, known as recovery agents, as specified by a recovery policy, described below. Similar to the FEK's encryption using the user's public key, the public portion of each recovery key pair is used to encrypt the FEK 60 using a data recovery field generator 80, (employing, for example, a suitable encryption algorithm, such as RSA, which need not be the same algorithm used to encrypt the FEK 60 with the user's public key). This list of encrypted FEKs is similarly stored along with the file 70 in a special EFS attribute called the Data Recovery Field (DRF) 82. Thus, only public portions of the recovery key pairs are needed for encryption of the FEK 60 in the DRF 82. Note that to facilitate proper operation, these public recovery keys are to be present at all times on an EFS system for normal file system 28 operations, since a user may wish to encrypt a file at any time. Recovery itself is expected to be a rare operation required only when users leave organizations, lose keys, and so on. As a result, recovery agents are also able to store the private portions 90 (FIG. 5) of the keys on smart cards, floppy disks, and/or other secure storage devices.

[0083] The EFS architecture is not limited to any particular encryption algorithm, but rather is fully algorithm agile and may use any cryptography algorithm for the various encryption phases. As a result, the EFS allows for the usage of better and better encryption algorithms as such algorithms advance technologically. Moreover, the user is able to choose from among available algorithms to select an algorithm having greater or less security (*i.e.*, based on how sensitive the user thinks the information is) versus the speed of encryption and decryption, (*i.e.*, more secure algorithms are generally slower). Thus, in the above description, DES is one such algorithm used to encrypt file data, while RSA is used to encrypt the FEK.

[0084] A choice of algorithms may be made available by providing the algorithm into an installable module separate from the file system and/or encrypting file system library (although, for example, a default algorithm may still be provided in the library). As generally represented in FIG. 2, a user can select an algorithm from an interchangeable (installable) cryptographic module 53 having a set of one or more suitable algorithms present therein. Alternatively, or in addition to, the algorithm set may be changed by replacing the interchangeable module with a different cryptographic module 53 containing a different algorithm set. For security, the interchangeable kernel mode cryptographic module 53 may be a kernel mode component, such as in the form of a single kernel mode export driver (a kernel-mode DLL). For example, the user (or an administrator) can choose a

given encryption/decryption algorithm for all files by default, on a per-file or per-directory basis, and so on. Once saved, information stored with the encrypted file can identify which algorithm was used to encrypt the data, whereby the appropriate algorithm for decrypting the data can be automatically selected for existing files. In any event, EFS is aware of the appropriate algorithm, and can notify the cryptographic module as to which one to use, *e.g.*, by calling a corresponding function of the module based on the algorithm or by passing an algorithm identifier to the cryptographic module.

[0085] By way of example, one such interchangeable kernel mode cryptographic module is a FIPS (Federal Information Processing Standards) system file, such as in the form of a single kernel mode export driver. The cryptographic boundary for this file is defined as the enclosure of the computer system on which the cryptographic module is to be executed.

[0086] To securely separate the interchangeable cryptographic module 53 from EFS library 47, the interchangeable cryptographic module 53 comprises a self-authenticating algorithm. The module 53 initializes before EFS, and does a self-check. When EFS initializes, the EFS library 47 calls the interchangeable cryptographic module 53 (driver) and receives a function table 53A in return, and stores it. For example, in an NTFS system, the table is acquired by building a function table request IRP (I/O request packet) and then sending the IRP to the interchangeable cryptographic module 53, which in turn returns the table.

Thereafter, when EFS performs encryption, EFS looks up the algorithm it will use in the function table, and uses it. Significantly, this provides a straightforward way for to change the algorithm and/or for an administrator or the like to replace an interchangeable cryptographic module 53 with an updated version. Such selection / replacement is independent and transparent to EFS, but allows EFS the flexibility to use different algorithms.

[0087] A preferred FIPS cryptographic module runs as a kernel mode export driver and encapsulates several different cryptographic algorithms in a cryptographic module that is accessible by other kernel mode drivers and can be linked into other kernel mode services (*e.g.*, to permit the use of FIPS 140-1 Level 1 compliant cryptography). The cryptographic module 53 may rely on the operating system for the authentication of users. The keys created within the cryptographic module 53 for one user are not accessible to any other user via the cryptographic module 53.

[0088] Once initialized, to use, for example, a DES or Triple DES function of the installable cryptographic module 53, a kernel mode system service provides a respective DES or Triple DES key. Keys are not stored by the module 53, but zeroed after the cryptographic module 53 completes a respective DES or Triple DES function with the keys.

[0089] In general, to encrypt or decrypt data, the EFS library 47 calls a respective function (*e.g.*, DES, 3DES) of the interchangeable cryptographic



module 53 with a pointer to an input buffer containing the data, and a pointer to an output buffer, a pointer to the key and a parameter specifying whether encryption or decryption is desired. Hashing functions also may be provided in the same module 53, along with key generation (*e.g.*, random keys), key entry and key output functions.

[0090] The following table summaries various states of a preferred FIPS interchangeable cryptographic module:

	Current State	Input	Output	Next State	Comment
1	Power Up	FIPS.SYS loads	NO_ERROR	Initialized	The Power Up state is entered when OS Loader calls the FIP.SYS driver entry point function DriverEntry() during system boot.
2	Power Up	FIPS.SYS not found	STATUS_UN_SUCCESSFUL	Init Error	(see comment for State 1 above)
2	Power Up	DES MAC check on cryptographic provider fails	STATUS_UN_SUCCESSFUL	Init Error	(see comment for State 1 above)
2	Power Up	One or more power-on cryptographic self-tests fail	STATUS_UN_SUCCESSFUL	Init Error	(see comment for State 1 above)
2	Power Up	System error	STATUS_UN_SUCCESSFUL	Init Error	(see comment for State 1 above)

3	Init Error	Automatic transition	No output	Power Down	The Init Error State is entered when FIPS.SYS's DriverEntry() fails as a result of either configuration errors ( <i>i.e.</i> not enough memory, etc.) or errors resulting from the power up self-tests.
4	Initialized	Key formatting operation ( <i>i.e.</i> FipsDesKey(), Fips3Des3Key() ) requested	No output	Key Initialized	The Initialized state is entered when FIPS.SYS's DriverEntry() returns successfully and the Windows Loader completes the loading of FIPS.SYS.
5	Initialized	Key formatting operation failure	Operation specific error message	Operation Error	
6	Operation Error	Automatic transition when keys have not yet been initialized	No output	Initialized	The Operation Error state is entered whenever an error occurs as a result of a cryptographic operation. FIPS.SYS will automatically transition back to either the Initialized or Key Initialized state depending on whether or not keys have been successfully formatted into a DESTable or DES3Table struct.
7	Key Initialized	Generic cryptographic operation failure	Operation specific error message	Operation Error	The Key Initialized state is entered after keys are formatted into a DESTable or DES3Table struct with FipsDesKey(), Fips3Des3Key() .
8	Operation Error	Automatic transition when keys have already been initialized	No output	Key Initialized	(see comment for State 6 above)
9	Key Initialized	Generic cryptographic operation ( <i>i.e.</i> FipsDes(), Fips3Des(), or FipsCBC ()) completed	NO_ERROR	Initialized	(see comment for State 7 above)

10	Initialized	Automatic transition when Windows XP Kernel calls the FIPS.SYS driver's unload function	NO_ERROR	Power Down	(see comment for States 4 and 5 above)
11	Power Down				The Power Down state is entered when OS calls the FIPS.SYS driver's unload function which was set in DriverUnload field of the DriverObject representing FIPS.SYS during the Power Up state.

## Data Decryption

[0091] **Fig. 4** conceptually shows the user decryption process. A user's private key 84 (or a recovery agent's private key 90) is used to decrypt the corresponding encrypted FEK item stored in the data decryption field 74. To accomplish the decryption of the key, each encrypted FEK item in the DDF 74 (and, if necessary the DRF 82) and the user's or recovery agent's private key are iteratively fed into an extraction mechanism 86 (in the EFS service 50), until a match is found which properly decrypts the FEK 60. From there, the FEK 60 is fed into a file decryption mechanism 88 which uses an appropriately corresponding decryption algorithm in a known manner to decrypt the encrypted text 68 into plain text 64. Note that when multiple decryption algorithms are available, the data decryption field may store information about the encryption algorithm so that the file decryption mechanism 88 uses the correct decryption algorithm. Only one

encryption algorithm can be used per file, although each of several files can have its own such algorithm.

[0092] While a file is open, the decrypted FEK 60 is saved by the file system 28 in association with the file 70. As shown in FIG. 6, with NTFS 28, stream control blocks 94 maintain file information for each open file, and each stream control block (*e.g.*, 941) corresponding to an encrypted file has a key context (*e.g.*, 961) pointed to thereby. The key context 96 maintains the information necessary to encrypt and decrypt a file during writes and reads to the disk, respectively. As described in more detail below, the FEK 60 is used to decrypt file data reads on a block-by-block basis, *i.e.*, random access to a large file will decrypt only the specific blocks read from disk for that file. The entire file is not necessarily decrypted.

## File Recovery

[0093] **Fig. 5** conceptually illustrates the recovery process. The recovery process is similar to user decryption, except that the process uses a recovery agent's private key 90 to decrypt the FEK 60 in the DRF 82. Consequently, no match will be found in the DDF 74 and thus the search for a match will continue into the DRF 82. To initiate a recovery, the recovery agent submits his or her private key 90, and a data recovery field extraction mechanism 92 (which is preferably the same data decryption field extraction mechanism 86 of FIG. 4 described above) iteratively

uses the agent's private key 90 to search the DDF 74. Then, since it is a recovery agent and not a user, no match is found in the DDF 74 and thus the extraction mechanism continues to search the DRF 82.

[0094] Regardless of whether it is a normal user opening or an opening for a recovery, once a key is found that is in the current context, (CryptoAPI 58 maintains a set of keys for each user), then the key is verified by comparing it to known information decrypted with that key. More particularly, at the time of encryption, the user's public key is appended to the FEK of the file, which is then encrypted with the FEK 60 as known information. If the found key decrypts the stored information to equal the known information, then the key is verified. This scheme provides a strong encryption technology as it provides one of many possible recovery agents with the ability to recover the file, thereby providing organizations with redundancy and flexibility in implementing recovery procedures.

[0095] EFS thus provides a built-in data recovery support, referred to as the "Recovery Policy". The preferred system enforces configuration of recovery keys, and is intentionally limited to only being usable when the system is configured with one or more recovery keys. The file recovery operation only divulges the randomly generated file encryption key 60, and not the user's or any other recovery agent's private key. As can be appreciated, this is ideal for most business

environments where an organization may need to recover data that was encrypted by an employee after an employee leaves the organization or loses his or her key.

[0096] The recovery policy, also known as the EFS policy, may be defined at the domain controller of an Active Directory Domain, whereby the policy is enforced at all machines in that domain. The policy contains the public keys of the recovery agents. As a result, the recovery policy is only under the control of the domain administrators, thereby providing controls on who can recover the data. To enable the use of encryption features on a standalone Windows 2000 or XP workstation in a home environment, as an added feature, EFS automatically generates recovery keys and saves them as machine keys, thereby reducing the administrative overhead for an average user.

[0097] In a domain, the recovery policy is sent to each of the machines on the domain, whereby even when not connected, the local machine maintains the policy therewith in a local security authority, LSA. In this manner, EFS can operate to encrypt files when not connected. Each time a machine joins the domain, or if the recovery policy changes, the policy is propagated to the machines in the domain when they are connected. Moreover, every time a file is opened, the metadata for that file is compared against the recovery policy to see if the recovery policy has changed, and if so, the metadata is updated with the FEK encrypted with the new user and/or recovery agent public key information. Safeguards, including

hashes (MD5), are used to ensure that the recovery policy is not changed, such as by a malicious user.

[0098] For example, a malicious user may wish to make a file unrecoverable by changing the DRF 82. To detect this, using a section of the DRF 82, a cryptography hash (MD5) is created, signed with the FEK 60 of the file, and stored with the DRF 82. Later, when the file is opened and the FEK 60 is obtained, the section is decrypted with the FEK 60 to see if the information stored in the DRF 82 matches. If so, the stored recovery policy is proper, otherwise the recovery policy is replaced with the current recovery policy of the domain.

[0099] If a machine is not part of a domain, it may have a local recovery policy, but the recover keys are generated by and kept on the machine as described above. In this manner, every file encrypted under EFS always has some recovery policy. If the machine later becomes part of a domain, the local recovery policy is wiped out and replaced with the domain policy.

## General Operation

[00100] Turning to **Fig. 10**, when an application 30 wishes to create or open an encrypted file, the application 30 first calls an appropriate API 32 requesting a new file be created in an encrypted directory, or an encrypted file be opened.

[00101] As shown in Fig. 10, once a create or open request is received, the I/O subsystem 56 arranges for passing the request as an IRP to the appropriate file

system, *e.g.*, NTFS 28, at steps 1000 - 1002. However, as described above, the IRP is first received by the EFS linked library 47 at step 1100 of FIG. 11, which recognizes the IRP as specifying an open/create operation.

[00102] The EFS linked library 47 begins the EFSCreateFile operation by performing some preprocessing as shown in **Fig. 11**. As represented by step 1102 of FIG. 11, the EFS linked library 47 allocates an EFS Context Block 981 for this file, and adds it to the EFS context block chain 98 (FIG. 7). Note that the EFS context block 981 is created for each new file, even though the file may not be encrypted, because the EFS linked library 47 does not know at this time whether the file is already encrypted or is to be encrypted. The EFS context block 981 includes status information, which is initialized with “No processing needed,” an IRP pointer pointing to the current file object, and an EFS metadata stream initialized to NULL. Lastly, at step 1104, the IRP is passed to NTFS 28.

[00103] As shown in step 1006 of FIG. 10, when NTFS 28 receives the IRP from the EFS linked library 47 and recognizes it as an NTFS 28 Create packet, NTFS 28 handles the create/open IRP. FIG. 12 generally shows how NTFS 28 handles the IRP. First, as represented by step 1200, information in the IRP is tested to determine if an existing stream is to be opened, or a new stream on file/directory is to be created. If an existing stream is to be opened, NTFS 28 opens the stream at step 1202. At step 1204, NTFS 28 determines if the file or its directory is



encrypted, and if so, at step 1208 calls the FSRTL 48 open /create callout, described below with reference to FIG. 13.

[00104] Conversely, if a new stream was created as determined by step 1200, NTFS 28 next determines at step 1206 if the parent directory is encrypted. If the parent is encrypted, at step 1208, NTFS 28 calls the FSRTL 48 open/create callout, as described below. Note that if neither the parent is determined to be encrypted (step 1206) nor the file or directory encrypted (step 1204), NTFS 28 does not make the FSRTL 48 callout, whereby NTFS 28 simply performs any tasks it needs to at step 1210 before returning to the EFS linked library 47.

[00105] The steps of the EFS create/open FSRTL 48 callout are generally represented in **Fig. 13**, wherein the FSRTL 48 callout begins by first examining the type of access requested by the user (step 1300). If an existing stream on the file or directory is opened without read, write, append or execute (R/W/A/E) access, the call is simply succeeded, as no encryption/decryption is needed, *e.g.*, the user wants to read attributes, and attributes are not encrypted. Otherwise, at step 1302, the FSRTL 48 searches the EFS context chain 98 for the appropriate file object corresponding to this file (allocated at step 1102 of Fig. 11). The Create/Open callout performs operations based on the type of file/directory, as set forth below.

[00106] If the type of file is an existing file (step 1304), then FSRTL 48 was called because either a new stream was created or an existing stream was opened. If so, the user needs to be verified, and the callout process continues to step 1400

of **Fig. 14**. At step 1400, the EFS metadata from the file is read using an (NtOfs) API. Then, at step 1402, the metadata that was read is set up in the context block 981 and the status on the block changed to indicate “User Verification Required.” Then the key context 96 is checked at step 1404, and if the NTFS key context 96 is NULL, then a key is needed. If the key context 96 is NULL (step 1404), this file was not read, and thus there is no possibility that decrypted file data is present in the cache memory. As a result, the context block is set to indicate “No Cache Check Needed” at step 1406. Lastly, if a new stream was created as determined by step 1408, the context block 981 is set to indicate “Turn On Encryption Bit” at step 1410.

[00107] If instead the type of file is a new file (step 1306 of FIG. 13), a new FEK and EFS metadata are needed. First, at step 1308, the EFS metadata is read from the parent directory using NtOfs API. Step 1310 sets up the metadata that was just read in the context block 981, and changes the status on the block to “New File Efs Required,” “Turn On The Encryption Bit” (step 1312) and “No Cache Check Needed (step 1314).”

[00108] If instead the file object type indicates a new directory (step 1320), only new EFS metadata is needed. There is no FEK in this case, because at present, streams in the directory are not encrypted. Accordingly, the EFS metadata from the parent directory is read at step 1322 (using NtOfs API). Then, at step 1324, the metadata that was just read is set up in the context block 981 and the

status on the block changed to “New Directory Efs Required,” “No Cache Check Needed” (step 1326) and “Turn On The Encryption Bit” (step 1328).

[00109] Lastly, if the type represents an existing directory (step 1332), either a new stream was created or an existing stream was opened. At present, no action is taken because the directory data streams are not encrypted. However, it can be readily appreciated that directory streams also may be encrypted in the same manner that file data streams are encrypted using the encrypting file system of the present invention.

[00110] As shown in **Fig. 12**, step 1210, the callout returns to NTFS 28, whereby NTFS 28 can perform any of its own operations. The file/create process returns to step 1010 (FIG. 10) the EFS filter EFS linked library 47 for post-processing.

[00111] The EFS Create/Open File post-processing process is represented in **Figs. 15-19**. Beginning at step 1500 of Fig. 15, the context block 981 is evaluated for “No Cache Check Required” status. If a cache check is required, the process branches to step 1502 where the caller’s security context along with the EFS ID for the file stored in the EFS stream are used by the EFS cache to check if this file was successfully opened by the user the recent past. If so, the call is succeeded since the cache already contains the appropriate information.

[00112] If not in the cache, step 1504 checks if read data, write data, append data or execute access is requested. If none of these are requested, but a new

stream was created as determined by step 1506, and the context block 981 indicates “Turn On Encryption Bit” status (STEP 1508), the EFS data stream is not needed and is released. Only the encryption bit on the stream needs to be turned on, which is performed at step 1510. The post processing is complete and the overall process returns to step 1012 of FIG. 10.

[00113] However, if none of the situations identified above with respect to FIG. 15 are satisfied, different operations need to be performed based on the status information, which is tested beginning at step 1600 of FIG. 16. First, if the status in the context block 901 indicates that user verification is required at step 1600, the EFS linked library 47 impersonates the security context (provided in the IRP) at step 1602, and at step 1604 calls the EFS service 50, passing the EFS metadata to request the FEK.

[00114] At step 1606, the EFS service 50 responds to the call by impersonating the context, and using information in the EFS metadata, looks up the user’s private key to decrypt the FEK. The EFS service 50 may also update the EFS metadata (step 1610) if the user’s key has been updated or any recovery keys are updated as determined by step 1608. In any event, at step 1612, the EFS service 50 verifies the integrity of the FEK 60 and returns all information back to the EFS linked library 47. More particularly, to verify integrity, a key integrity block is constructed as follows:

[00115]  $[F(\text{FEK}, P_{uk}), \text{FEK}]P_{uk}$

[00116]           where

[00117]  $F()$  is a suitable hash function (preferably MD5),

[00118]  $P_{uk}$  is the user's public key, and

[00119]  $[]P_{uk}$  denotes encryption with the user's public key

[00120] Consequently, when it is believed that a valid FEK has been decrypted with a user's public key, the block above may be computed with the present information and compared to the block stored on the file using a checksum or hash to verify integrity. If they match, the key integrity is verified.

[00121] Alternatively, if the status did not indicate that user verification was required (step 1600), but instead indicated that a new file FEK was required, step 1614 branches to step 1700 of FIG. 17. At step 1700, the EFS linked library 47 impersonates the securing context in the IRP, and at step 1702 calls the EFS service 50 passing the parent directory's EFS metadata, requesting a new FEK and EFS metadata. At step 1704, the EFS service 50 impersonates the context and generates a random FEK 60. If the user does not have a key as determined by step 1706, at step 1708, the EFS service 50 may auto-generates a key pair for the user or automatically request a certificate and key pair from a certificate authority. Lastly, step 1710 creates the EFS metadata stream with the FEK 60 encrypted under the user's public key and the recovery agent's public keys. The EFS service 50 also encrypts the FEK 60 using all the public keys in the parent directory's EFS

metadata so that users who are allowed access to the parent directory also have access to the file (provided NTFS 28 access control lists allow such access).

[00122] If neither step 1600 nor 1614 was satisfied, the post-process branches to step 1616 to determine if the EFS context indicated that a new directory FEK is required. If so, the post-process branches to step 1800 of FIG. 18 wherein the EFS linked library 47 impersonates the securing context in the IRP. Step 1802 calls the EFS service 50, passing the parent directory's EFS metadata and requesting new EFs metadata. Note that no FEK is needed, as directory streams are not encrypted at this time. However, in the future, directory streams may also be encrypted in the same manner that file streams are encrypted in accordance with the present invention.

[00123] In any event, at step 1804, the EFS service 50 impersonates the context, and, using an empty FEK, creates the EFS metadata stream. Step 1706 checks to see if the user does not have a key, and if not, at step 1708, the EFS service 50 auto-generates a key pair for the user. Then, at step 1810, the empty FEK is encrypted under the user's public key and the recovery agent's public keys, and the FEK is also encrypted using all the public keys in the parent directory's EFS metadata so that users allowed access to the parent directory also have access to the file if the access control lists allow such access.

[00124] Ultimately, regardless of which of the three statuses were in the context, the post-process winds up at step 1900 of FIG. 19 to issue an appropriate

FSCTL. Two such FSCTL calls are available, FSCTL\_SET\_ENCRYPTION and FSCTL\_ENCRYPTION\_FSCTL\_IO. FSCTL\_SET\_ENCRYPTION tells NTFS 28 to turn on or turn off the encryption bit for a stream. The FSCTL\_ENCRYPTION\_FSCTL\_IO is a miscellaneous FSCTL used for performing a number of operations, described below.

[00125] To this end, the FSCTLs are accompanied by a data structure 100, as shown in FIGS. 8 and 9. The data structure includes a public code so that NTFS 28 can differentiate between the two types of FSCTL calls, along with an EFS subcode to more particularly define the operation for the EFS linked library 47 and/or the EFS service 50. The data structure also includes EFS data specifying either FEK information (FIG. 8) or file handle information (FIG. 9), and, at times, EFS metadata. For purposes of security, the EFS subcode and EFS data fields are encrypted with the session key established when the EFS service 50 is initialized, as described above.

[00126] In use, the FSCTL\_SET\_ENCRYPTION may be issued, for example, to turn on the encryption bit for a file when that file is first put into an encrypted directory. The subcode indicates whether the encryption bit should be turned on or off. For such an operation, the FEK 60 is already known, and thus as shown in FIG. 8, the EFS data includes the FEK, and the FEK encrypted with the session key. Note that all but the public code is encrypted with the session code. To verify the integrity and the source of the data structure 100, the encrypted

portion of the data structure is decrypted. Then, the encrypted FEK is decrypted and compared with the other FEK, and if equal, the structure is verified. The EFS stream, if available may also be compared with the EFS metadata, if otherwise known. Since the FEK is not always known, a similar verification is performed using the session key and the file handle as shown in FIG. 9. A repeated file handle is actually used in the appropriate fields (FIG. 9) so as to equal a length of eight bytes.

[00127] The callouts to the FSRTL 48 (via FileSystemControl\_1 or FileSystemControl\_2, and passed through NTFS 28) are used to overwrite attributes or set attributes depending on an accompanying subcode. Two bits of the subcode represent the overwrite attributes or set attributes operations. Note that when the EFS stream is to be written for a new file, FILE\_SYSTEM\_CONTROL\_1 is used with the FSRTL 48 callout to also turn on the encryption bit (as performed by NTFS 28, described above). Alternatively, FILE\_SYSTEM\_CONTROL\_2 is used with the callout when no change to the encryption bit is needed, for example, if the user has simply changed user keys.

[00128] Regardless which is used, one bit of the subcode represents the operation “Set EFS KeyBlob,” which indicates to the FSRTL 48 that new encryption key information is available and needs to be entered into the appropriate key context. Another bit represents the operation “Write EFS Stream.” Write EFS Stream is issued by the EFS service 50, such as when the user or



recovery agent has changed a public key and the file metadata needs to be rewritten with the EFS metadata in the data structure 100. One other subcode represents “Get EFS stream,” which results in the current EFS attributes for a file being written into the EFS field, such as when a user wants to export a stream or wants to know a key name.

[00129] Thus, returning to FIG. 19, step 1900 tests if the only requirement is to turn on the encryption bit, and if so, issues the FSCTL\_SET\_ENCRYPTION control at step 1902 with the subcode indicating that the bit should be turned on. Of course, the other EFS data including the session key, handle, handle and encrypted copy of same is also in the data structure 100 for verification purposes. In any event, the FSCTL reaches NTFS 28, which turns around and sets the encryption bit on the stream (and on the file if not already on) and calls an FSRTL callout to pass the information to the FSRTL 48.

[00130] If step 1900 is not satisfied, the ENCRYPTION\_FSCTL\_IO FSCTL needs to be called with an appropriate subcode in the data structure 100. Thus, if the status is “User Verification Required” (step 1904), the subcode is set to KeyBlob at step 1905. Next, if the metadata has been updated as determined by step 1906, step 1907 sets the Write EFS Stream subcode bit before the call at step 1914. Otherwise, if the status is “New File FEK Required” (step 1908), the subcode is set to KeyBlob and Write EFS Stream at step 1910, *i.e.*, both bits are set. If neither of these, then the status is “New Directory FEK Required,” and the

subcode is set to Write EFS Stream at step 1912, *i.e.*, only the other bit is set. The FSCTL is issued at step 1914.

## Read and Write

[00131] Turning to **Fig. 20**, when an application 30 requests to read some data from the open file (step 2000), the I/O subsystem 56 receives the read request and passes it as an IRP to the appropriate file system 28, *e.g.*, NTFS. First, however, the IRP is received by the EFS linked library 47, which recognizes the IRP as corresponding to a read request, and as a result, directly hands the IRP to NTFS 28 at step 2002. At step 2004, NTFS 28 reads the encrypted data from disk into a buffer just as it would read the plaintext data for any other file. However, for the encrypted file, the file system 28 recognizes at step 2006 that this file is encrypted, and at step 2008 remembers and gets the key context 961 that the encryption EFS linked library 47 earlier had returned from the create/open callback. At step 2010, NTFS 28 uses the AfterReadProcess callback and provides the registered function with the data and enough information, including the encryption context, to decrypt the data. In general, the information includes the offset into the file, a pointer to the read buffer, the length to read, and the key. At step 2012, the encryption EFS linked library 47 decrypts the data and returns it to the file system 28, whereby at step 2014 the file system 28 then returns this plaintext through the I/O subsystem 56 to the application in the normal way. Note

that however that certain NTFS 28 internal metadata streams containing file indexing and other such information are not encrypted. NTFS 28 recognizes these streams at step 2006, and temporarily skips over steps 2008 - 2012 for these particular streams.

[00132] As shown in **Fig. 21**, when an application 30 requests to write data to the open file (step 2100), the I/O subsystem 56 receives the write request and passes it as an IRP to the appropriate file system 28, *e.g.*, NTFS. First, however, the IRP is received by the EFS linked library 47, which recognizes the IRP as corresponding to a write request, and as a result, directly hands the IRP to NTFS 28 at step 2102. At step 2104, NTFS 28 copies the write data into a separate buffer so that no changes can be made to the data that is to be written. For the encrypted file, the file system 28 recognizes at step 2106 that this file is encrypted, and at step 2108 remembers and gets the key context 961 that the encryption EFS linked library 47 earlier had returned from the create/open callback. At step 2110, NTFS 28 uses the BeforeWriteProcess callback and provides the function with the data and enough information, including the encryption context, to encrypt the data. At step 2112, the encryption EFS linked library 47 encrypts the data and returns it to NTFS 28, whereby at step 2114 the file system 28 then writes the now-encrypted data in the separate buffer to the non-volatile storage 40 in the normal way, *i.e.*, as if it was plaintext data for any other file. Again note that the NTFS 28 internal metadata streams containing the file indexing and other such information are not to

be encrypted. NTFS 28 recognizes these streams at step 2106, and temporarily skips over steps 2108 - 2112 for these particular streams.

## Encrypt and Decrypt File APIs

[00133] EFS also provides APIS 32 to facilitate encryption and decryption of stored files. The Win32 EncryptFile API is used to encrypt a plaintext file/directory. As shown in **Fig. 22**, with this API, the application 30 (user) provides the name of the file to encrypt at step 2200, and this call translates into a call to the EFS service 50 to do the operation. At step 2202, the EFS service 50 opens the file on the user's behalf, makes a backup copy for crash recovery purposes and at step 2204 marks it for encryption by issuing the SET\_ENCRYPT file control (FSCTL). At step 2206, the EFS service 50 then reads data from each stream in the copy and writes the data back to the original file at step 2208. Note that during the write operation, because the encryption bit is set, the data is automatically encrypted before being written to the disk. The process is repeated via step 2210 until all data streams are written. If this call completes successfully (step 2212), the backup is deleted at step 2214, otherwise the original file is restored at step 2216 and the call is failed. In the case of a directory, the directory is simply marked encrypted, as there is no data to encrypt. Note that, as described above, NTFS 28 internal metadata streams are not encrypted.

[00134] A WIN32 DecryptFile API is also provided by the EFS service 50, and is the converse operation of the encrypt file/directory operation. As shown in FIG. 23, at steps 2300 –2302, the EFS service 50 is provided with the file name and opens the file on the user's behalf. Steps 2306 - 2310 read the data from all streams and write those streams into a copy, which is plaintext, as decryption happens transparently. At step 2312, the service then issues the decrypt file control to delete the metadata and remove the encryption attribute. Then, as shown by steps 2314 – 2318, the API writes back all the data streams from the copy over the original, which are written in plaintext. If this completes successfully, the copy is deleted at step 2322, otherwise the original is restored at step 2324. In the case of a directory, the directory is simply marked as decrypted to delete the metadata and the attribute, as there is no data to decrypt.

[00135] As can be appreciated, EFS file encryption is supported on a per file or entire directory basis (although NTFS 28 operates per stream). Directory encryption is transparently enforced, *i.e.*, all files (or subdirectories) created in a directory marked for encryption are automatically created encrypted. Moreover, file encryption keys are per file, making them safe against move/copy operations on the file system volume. Unlike existing application-level schemes, the file need not be decrypted before use, since, as will become apparent below, the encryption and decryption operations will be done transparently and on the fly when bytes travel to and from the disk. EFS will automatically detect the encrypted file and

locate the user's key from a key store. The mechanisms of key storage are leveraged from CryptoAPI, and as a result the users will have the flexibility of storing keys on secure devices such as smart cards and/or floppy disks.

[00136] Moreover, EFS cooperates with the underlying file system 28 (*e.g.*, NTFS), whereby EFS facilitates the writing (by a properly developed application program) of encrypted temporary files. With such an application program, when temporary files are created, the attributes from the original file are copied to the temporary file making the temporary copy also encrypted. In addition, the EFS linked library 47 is a Windows NT kernel mode driver, which uses the non-paged pool to store file encryption key, thereby ensuring that the key never makes it to the page file.

[00137] The EFS architecture allows file sharing between any number of people by simple use of the public keys of those people. Each user can then independently decrypt the file using their private keys. Users can be easily added (if they have a configured public key pair) or removed from the clique of sharers.

[00138] In a stand-alone configuration, EFS allows users to start encrypting/decrypting files with no administrative effort to set up a key pair and certificate, *i.e.*, EFS supports auto-generation of a key for the user if one is not configured. With a domain configuration, an administrator may set up a domain policy for all users and machines. Lastly, EFS will also support encryption/decryption on remote files stored on file servers.

## Client-Based Encryption

[00139] In alternate embodiments, an EFS system is adapted to enable client computers to manage encryption and decryption of files in a client-server computer network, or between peer computers in a peer-to-peer network. Users at a client computer can create, open, read from, or write to an encrypted file located on a remote computer. In addition, users can encrypt or decrypt a file stored on a remote computer, and can grant or deny users access to an encrypted file. This permits users to retain control of the encryption state of their files.

[00140] **Fig. 24** is a schematic illustration of components of an exemplary EFS architecture adapted to permit client-based encryption and decryption. Referring to Fig. 24, a client-side computing device comprises a local security authority (LSA) 2410 that operates in the user mode of the operating system and a redirector 2415 that operates in the kernel mode of the operating system. The redirector 2415 links to the EFS, *e.g.*, through APIs. A server side computing device comprises an LSA 2425 that operates in the user mode of the operating system and a server module 2430 that operates in the kernel mode of the operating system. Server module 2430 interacts with the NTFS 2435, which in turn links to the EFS 2440, *e.g.*, through APIs. Server module 2430 may also link to the EFS 2432, *e.g.*, through APIs.

[00141] Client computing device and server computing device may be implemented in a network such as, *e.g.*, a local area network (LAN), a wide area network (WAN) or another suitable network. It will be appreciated that the network may comprise multiple clients and multiple servers, some or all of which may be configured as depicted in Fig. 24. Clients and servers in a network may share files, typically pursuant to a file sharing protocol. In the WINDOWS® NT operating system, the client redirector 2415 and the server module 2430 are configured to support the Server Message Block (SMB) protocol for compatibility with existing MS-NET and LAN Manager servers (thus allowing access to MS-DOS, WINDOWS, and OS/2 systems), as well as the NCP protocol for communication with Netware-compatible servers and clients. In alternate operating systems, the redirector 2415 and the server module 2430 may be configured to support other file sharing protocols, *e.g.*, the network file sharing (NFS) protocol.

#### Initialization

[00142] In operation, a redirector needs the process ID (PID) of the LSA to decrypt a FEK, and needs to initialize the FIPS driver. Thus, the redirector's initialization process is dependent on the LSA's EFS module being loaded. In an exemplary embodiment the client implements a routine to synchronize initialization of the redirector with the EFS module of the LSA. On initialization,



the redirector 2415 attempts to create an event (EFSSmbInitEvent). If the event is already created, then the redirector 2415 knows that the EFS module of the LSA is already loaded. If the event is not created, then the redirector spawns a new thread to wait until the event is set.

[00143] Similarly, when the EFS module of the LSA finishes loading, it attempts to create the EFSSmbInitEvent. If it fails, *i.e.*, if the redirector 2415 created the EFSSmbInitEvent event, then the EFS module of the LSA opens the event and signals that it is active. This procedure enables the redirector 2415 to synchronize its initialization with the EFS module of the LSA.

[00144] Similarly, the server module 2430 needs the EFS driver session key, which is generated in the server's LSA, to communicate with the EFS, and to acquire the \$EFS stream. Therefore, in an exemplary embodiment the server implements a similar routine to synchronize initialization of the server module 2430 with the EFS module of the LSA. On initialization, the server module 2430 attempts to create an event (EFSSrvInitEvent). If the event is already created, then the server module 2430 knows that the EFS module of the LSA is already loaded. If the event is not created, then the redirector spawns a new thread to wait until the event is set.

[00145] When the EFS module of the LSA finishes loading, it attempts to create the EFSSrvInitEvent. If it fails, *i.e.*, if the server module 2430 created the EFSSrvInitEvent event, then the EFS module of the LSA opens the event and

signals that it is active. This procedure enables the server module 2430 to synchronize its initialization with the EFS module of the LSA.

#### Operations on the Remote \$EFS Stream

[00146] Implementing client-based encryption and decryption requires the ability to read and modify the \$EFS stream of the remote file over the SMB file transfer protocol. In an exemplary embodiment, remote metadata (*i.e.*, the \$EFS stream) is acquired by implementing an SMB message (NT\_TRANSACT\_EFS\_FSCTL). A data structure associated with the SMB message contains a control code indicating an operation for the server to perform, *e.g.*, obtaining the EFS metadata (SMB\_EFS\_FSCTL\_GET\_FILE\_METADATA). Read access to remote metadata for a particular file may be granted to any user who has READ\_ATTRIBUTES for the file.

[00147] To get the \$EFS metadata from the SMB context, the server invokes a FSCTL\_ENCRYPTION\_FSCTL\_IO, passing EFS\_GET\_ATTRIBUTE as the control code. The server encrypts parameters passed to this FSCTL with the session key initialized at startup.

[00148] Modification of remote \$EFS metadata may be implemented by an SMB message. A data structure associated with the SMB message contains a control code indicating an operation for the server to perform, *e.g.*, modifying the EFS metadata (SMB\_EFS\_FSCTL\_SET\_FILE\_METADATA). The modification

is implemented in two stages. First, it is determined whether a user at the client has appropriate access to make the requested modification. If the user has WRITE\_DATA access, then write operations on the \$EFS stream are enabled. By contrast, if a user does not have WRITE\_DATA access, but has a key which can decrypt the FEK and WRITE\_DATA\_ATTRIBUTE permission, then the user will be allowed to modify their DDF. Proof of possession of the decryption key provides for allowing only authorized users to modify the files. Second, once the access check has been performed, writing the \$EFS stream may be accomplished by calling FSCTL\_ENCRYPTION\_FSCTL\_IO, and specifying EFS\_OVERWRITE\_ATTRIBUTE as the control code.

[00149] In an exemplary embodiment, the FSCTLs may be encapsulated in a remote file context that can be used for encrypting and decrypting data. The redirector passes the context to encryption and decryption routines as a handle. The following exemplary algorithm may be used to demonstrate that the user has a key with which the user can decrypt the FEK. The client obtains the \$EFS stream of the file and updates the \$EFS stream before sending it back to the server. By way of example, in a rekey operation the client may update its DDF entry in the \$EFS stream before sending it back to the server. The client sends the following information to the server: the \$EFS stream; the EFS certificate associated with the Smartcard private key; and the RSA signature value, (which may be computed over the SHA1 hash of the \$EFS stream). The server computes the hash of the EFS

certificate and tries to find a match in a DDF entry in the \$EFS stream of the server copy of the file. If a match is found, the server then verifies the signature on the certificate and verifies the signature value using the hash of the \$EFS stream. No certificate chain is built and no chain validation is performed. If the signature can be verified, the client has proven possession of the private key and the server will update the server \$EFS stream with the client copy.

### Create File Operations

[00150] **Fig. 32** is a flow diagram illustrating exemplary operations for creating a file. Creating a new encrypted file involves generating a \$EFS stream with the FEK. When encryption is performed on the client, the \$EFS may be generated on the client so that the FEK need not be available on the server. The \$EFS may be transmitted in either an encrypted form or an unencrypted form from the client to the server and set on the file on the server.

[00151] Referring to Fig. 32, at operation 3210 an application requests a file creation. At operation 3214 the redirector sends a request to the SMB server. The server receives the request, and at operation 3218 the server sends a request to the NTFS. The server knows that an EFS file is to be created, and the request to the NTFS includes a flag indicating that the file is to be an EFS file. At operation 3222 the server sends a request for the \$EFS to the redirector. The redirector receives the request and, at operation 3226, the redirector calls the EFS library to

get the \$EFS. The EFS library calls EFS user mode code to create the \$EFS (operation 3230) and at operation 3234 the redirector sends the \$EFS to the server. The server receives the \$EFS and, at operation 3238, the server sends a request to the NTFS to write the \$EFS. At operation 3242 the NTFS calls the EFS library on the server to write the \$EFS. At operation 3246 the server returns a handle to the client, and at operation 3250 the client returns the handle to the application.

### Read and Write Operations

[00152] **Fig. 25** is a flow diagram illustrating an exemplary read operation. At operation 2510 an application on a client requests to read data from an encrypted file resident on a server. The request is sent to the I/O system, and the redirector 2415 intercepts the request, at operation 2512. The read request may be implemented via an API that accepts the following arguments: (1) FileHdl - a handle to the file from which to read data; (2) EfsFileContext - context information needed to decrypt the file (which is returned from the EfsSmbGetFileContext API); StartingFileOffset - the starting offset in the file from which to begin the read; (4) ReadDataBuffer - a buffer in which to store the resulting read; and ReadDataSize - the amount of data to read. In an exemplary embodiment the CreateFile path is configured not to access the user private keys on the server if the read request comes from a CSE-aware component. A pointer to a flag is passed to the EFS on file open or file create. If the flag is set, then the EFS returns to the

caller using a dummy encryption context, which performs an identity transform on the encrypted data.. Thus, a file handle can be opened for read and write operations without requiring a private key to be present on the server. When a file is created, a pointer to the \$EFS stream is passed with the flag.

[00153] The redirector 2415 passes the read request to the server module 2430 at operation 2514, which performs a raw read of the encrypted data, at operation 2516. In an exemplary embodiment, the server module 2430 invokes the NTFS 2435 to perform a raw read, *e.g.*, by issuing an FSCTL, and the EFS library breaks the encrypted data in to packets of 64KB for transfer.

[00154] At operation 2518 the server module 2430 returns the data to the redirector 2415, and at operation 2520 the redirector 2415 validates the received data. In an exemplary embodiment the encrypted data is returned in an ENCRYPTED\_DATA\_INFO structure that includes a pointer to the encrypted data. The pointer is retrieved and the redirector validates that the encrypted data is within the valid data length parameters of the ENCRYPTED\_DATA\_INFO structure. If the encrypted data is not within the valid data length parameters of the ENCRYPTED\_DATA\_INFO structure, then data beyond the 512 byte boundary past the valid data length is discarded.

[00155] At operation 2522 the redirector 2415 invokes the EFS to decrypt the encrypted data. Decryption is described in greater detail below. At operation 2524

the redirector 2415 returns the decrypted data (or a pointer to the data) to the calling application.

[00156] In addition to read/write operations, the redirector and EFS also implement a limited subset of NtSetInformationFile functionality. Specifically, the redirector and EFS simulates FileValidDataLengthInformation operations that require possession of the FEK. No special processing is required if the valid data length (VDL) is being reduced. If the VDL is being increased, then the client writes zeros from the original VDL boundary, (or from the next-highest 512 byte boundary) to the new VDL (or to the next-highest 512 byte boundary). The VDL may not be aligned upward within the server. Instead, the client may use the BytesWithinValidDataLength field of the ENCRYPTED\_DATA\_INFO structure to control this value.

[00157] **Fig. 26** is a flow diagram illustrating an exemplary write operation. At operation 2610 an application on a client requests to write data to an encrypted file resident on a server. At operation 2612 the redirector obtains the write request, e.g., from the I/O manager, and at operation 2614 the redirector forwards a read request to the server. The server receives the read request and, at operation 2616, the server reads the encrypted data. At operation 2618 the server forwards the encrypted data to the redirector.

[00158] At operation 2620 the redirector invokes the EFS to overwrite the data identified in the write request. The EFS client decrypts portions of the

encrypted data returned by the server, overwrites the data identified in the write request, and encrypts the modified data. At operation 2622 the redirector forwards a write request to the server using the modified, encrypted data generated in response to the write request. The write request may be implemented via an API that accepts the following arguments: (1) FileHdl - a handle to the file to which data will be written; (2) EfsFileContext - the context information needed to encrypt this file (this information is returned from EfsSmbGetFileContext); (3) StartingFileOffset - the starting offset in the file from which to begin the write; (4) a flag that indicates whether data is already block aligned and may be written directly to the server; (5) WriteDataSize - the amount of data to write; and (6) WriteDataBuffer – a buffer in which to store the resulting read.

[00159] The server module 2430 performs a raw read of the encrypted data to determine appropriate cluster-aligned offsets for the write operation. If necessary, the server module 2430 determines appropriate offsets for writing to the file. At operation 2624 the server module writes the encrypted data to the file, *e.g.*, by invoking the NTFS. At operation 2626 the server returns a response to the write request to the redirector.

[00160] In an exemplary implementation the client determines if the requested write is past the VDL. If not, then the regular write path may be used. By contrast, if the write is past the VDL, then the client sets the VDL to the



starting offset of the requested write, e.g., by writing zeros to the new 512 byte boundary.

## Encrypt and Decrypt File Operations

[00161] **Fig. 27** is a flow diagram illustrating an exemplary encrypt operation. At operation 2710 an application on the client requests an encryption operation on a first file stored at the server. The operating system forwards the encrypt request to the client LSA 2410, which forwards the encrypt request to the server 2430, *e.g.*, using a remote procedure call (RPC), at operation 2714. The server 2430 receives the encrypt request, and in response the server creates a temp file at operation 2716, *e.g.*, by calling a CreateFile operation from the NTFS using the template information from the first file. The temp file then contains the metadata from the first file. The server forwards the name of the temp file to the client at operation 2718.

[00162] At operation 2720 the client opens the first file and the temp file, and at operation 2722 the client writes an \$EFS stream to the temp file indicating that the temp file is an encrypted file. At operation 2726 the client writes the data from the first file to the temp file. The EFS encrypts the data during the write process. At operation 2728 the client replaces the first file with the temp file. In an exemplary embodiment, the client may simply close and delete the first file, then re-name the temp file to have the same name as the original file.

[00163] The operations described with reference to Fig. 27 enable client-based encryption of files resident on a remote server. The FEK need not be exposed to the server during the process.

[00164] **Fig. 28** is a flow diagram illustrating an exemplary decrypt operation. At operation 2810 an application on a client requests a decrypt operation on a first file. At operation 2812 the client obtains the \$EFS stream of the first file. In an exemplary embodiment the client obtains the \$EFS stream by issuing a remote FSCTL call, as described above. At operation 2814 the client extracts the FEK from the \$EFS stream. At operation 2816 the client forwards the FEK to the server with a request to decrypt the file.

[00165] At operation 2818 the server creates a temp file, and at operation 2820 the server writes the first file metadata from the first file to the temp file. At operation 2822 the server decrypts the first file using the FEK supplied from the client, *e.g.*, by invoking the EFS to implement a decryption routine. At operation 2824 the server writes the decrypted data to the temp file, and at operation 2826 the server replaces the first file with the temp file. In an exemplary embodiment, the server may simply close and delete the first file, then re-name the temp file to have the same name as the original file. In an alternate implementation, the server could return the encrypted text to the client, which could perform the decryption operation and any necessary writes back to server.

## Encrypted Raw File Operations

[00166] As described above, system interfaces are provided for opening, reading, and writing to an encrypted raw file. As used herein, the term “raw file” shall be construed to mean a file that is not transformed through decryption during a read process and retains its state and metadata. The OpenRawFile interface allows a user to open an encrypted file without file system read access, and without setting up a file encryption key to do transparent reads and writes. Similarly, a ReadRawFile interface allows a user to read all the data from the file, including the encryption metadata, as a contiguous opaque stream that can be backed up and later restored. A WriteRawFile interface allows a user to write all the data to the file from its backup, including the encryption metadata, to re-create the encrypted file. These interfaces may be leveraged to supply the capability to open, read, or write encrypted raw files in a client-server context or in a peer-to-peer context as well.

[00167] **Fig. 29** is a flow diagram illustrating an exemplary method for opening an encrypted raw file in a client-server context. At operation 2910 a request to open a raw encrypted file is received, *e.g.*, from an application or from another user interface. At operation 2915 it is determined whether the file path is a local path, *i.e.*, a path on the client, or a remote path, *i.e.*, a path on a server. If the file path is local, then control passes to operation 2920, and the OpenRawFile Interface is called, and the process ends at operation 2925.

[00168] By contrast, if the file path is remote, then control passes to operation 2930, and the OpenRawFile interface is called with the remote file path. If the user is a backup operator on the remote device, then the OpenRawFile interface call will execute successfully (operation 2935) on the remote device and the process ends at operation 2925.

[00169] If the user is not a backup operator on the remote device, the OpenRawFile interface call will fail (operation 2935), and control passes to operation 2940, where a remote file access flag is set in the file context structure passed with the OpenRawFile interface call. At operation 2945 the OpenRawFile interface is called with the remote file access set. In response to the remote access flag, the OpenRawFile interface suppresses any inquiry into file system or network share access levels, and the process ends at operation 2925.

[00170] A raw read and a raw write from an encrypted file in a client-server context may be performed in substantially the same manner as a local raw read. In an exemplary implementation, the redirector will use an SMB message to communicate with the remote or second machine.

### Adding Users to and Removing Users From Encrypted Files

[00171] **Fig. 30** is a flow diagram illustrating an exemplary procedure for adding a user to an encrypted file. Referring to Fig. 30, at operation 3010 the client requests the \$EFS stream from the server, *e.g.*, using a remote procedure call

as described above. In response to the remote procedure call, the server makes the \$EFS stream available to the client. At operation 3015 the client generates a data decryption field (DDF) for a user. At operation 3020 the client modifies the \$EFS stream to include the DDF of the added user. And at operation 3025 the client writes the \$EFS stream back to the server, *e.g.*, using a remote procedure call as described above.

[00172] **Fig. 31** is a flow diagram illustrating an exemplary procedure for deleting a user from an encrypted file. Referring to Fig. 31, at operation 3110 the client requests the \$EFS stream from the server, *e.g.*, using a remote procedure call as described above. In response to the remote procedure call, the server makes the \$EFS stream available to the client. At operation 3115 the client reads the \$EFS stream and locates the DDF in the \$EFS stream that corresponds to the user being removed. At operation 3120 the client removes the DDF from the \$EFS stream. And at operation 3125 the client writes the \$EFS stream back to the server, *e.g.*, using a remote procedure call as described above.

[00173] Thus, it will be apparent that the operations described in Figs. 24-32 permit a user at a client computer to control the encryption and/or decryption of files in a client-server network.

[00174] Although certain implementations have been described in language specific to structural features and/or methodological operations, it is to be understood that the subject matter defined in the appended claims is not

necessarily limited to the specific features or operations described. Rather, the specific features and operations are disclosed as preferred forms of implementing the claimed present subject matter.